

Best-Effort Strategies for Losing States

Technical Report

Marco Faella

Computer Science Division
Physics Department
Università di Napoli “Federico II”, Italy

Abstract. We consider games played on finite graphs, whose goal is to obtain a trace belonging to a given set of winning traces. We focus on those states from which Player 1 cannot force a win. We explore and compare several criteria for establishing what is the preferable behavior of Player 1 from those states.

Along the way, we prove several results of theoretical and practical interest, such as a characterization of admissible strategies, which also provides a simple algorithm for computing such strategies for various common goals, and the equivalence between the existence of positional winning strategies and the existence of positional subgame perfect strategies.

1 Introduction

Games played on finite graphs have been widely investigated in Computer Science, with applications including controller synthesis [PR89,ALW89,MPS95,dAFMR05], protocol verification [KR01,BBF07], logic and automata theory [EJ91,Zie98], and compositional software verification [dAH01].

In such games, we are given a finite graph, whose set of states is partitioned into Player 1 and Player 2 states, and a *goal*, which is a set of infinite sequences of states. The game consists in the two players taking turns at picking a successor state, giving rise to an ever increasing and eventually infinite sequence of states. A (deterministic) *strategy* for a player is a function that, given the current history of the game (a finite sequence of states), chooses the next state. A state s is said to be *winning* if there exists a strategy that guarantees victory regardless of the moves of the adversary, if the game starts in s . A state that is not winning is called *losing*.

The main algorithmic concern of the classical theory of these games is determining the set of winning states. In this paper, we shift the focus to losing states, since we claim that many applications would benefit from a theory of best-effort strategies which allowed Player 1 to play in a rational way even from losing states.

For instance, many game models correspond to real-world problems which are not really competitive: the game is just a tool which enables to distinguish internal from external non-determinism. In practice, the behavior of the adversary may turn out to be random, or even cooperative. A strategy of Player 1 which does not “give up”, but rather tries its best at winning, may in fact end up winning, even starting from states that are theoretically losing. For instance, such is the case in [dAFMR05], where games are used to model the interaction between the scheduler of an operating system and the applications being scheduled.

In other cases, the game is an over-approximation of reality, giving to Player 2 a wider set of capabilities (i.e., moves in the game) than what most adversaries actually have in practice. Again, a best-effort strategy for Player 1 can thus often lead to victory, even against an adversary which is strictly competitive.

In this paper, we consider and compare several alternative definitions of best-effort strategies. As a guideline for our investigation, we take the application domain of automated verification and synthesis of open systems. Such domain is characterized by the fact that, once good strategies for a game have been found, they are intended to be actually implemented in hardware or software. As a consequence, we tend to favor best-effort criteria that are as discriminating (i.e., specific) as possible while still giving rise to efficient strategies. While having such application domain in mind, we still put the main focus of the present work on theoretical issues, leaving to future investigations a discussion of what criterion is more suitable for any specific application.

Best-effort strategies. The classical definition of what a “good” strategy is states that a strategy is *winning* if it guarantees victory whenever the game is started in a winning state [Tho95]. This definition does not put any burden on a strategy if the game starts from a losing state. In other words, if the game starts from a losing state, all strategies are considered equivalent.

A first refinement of the classical definition is a slight modification of the game-theoretic notion of *subgame-perfect equilibrium* [OR94]. Cast in our framework, this notion states that a strategy is good if it enforces victory whenever the game history is such that victory can be enforced. We call such strategies *strongly winning*, to avoid confusion with the use of subgame (and subarena) which is common in computer science [Zie98]. It is easy to see that this definition captures the intuitive idea that a good strategy should “enforce victory whenever it can” better than the classical one.

Next, consider games where victory cannot be enforced at any point during the play. Take the Büchi game in Figure 1¹, whose goal is to visit infinitely often s_0 . No matter how many visits to s_0 Player 1 manages to make, he will never reach a point where he can enforce victory. Still, it is intuitively better for him to keep trying (i.e., move to s_1) rather than give up (i.e., move to s_2). To capture this intuition, we consider two further notions of best-effort strategy.




Fig. 1: A game where victory cannot be enforced.

One such notion which is well known in the literature is the one of optimal strategy in a particular type of game called a *Markov Decision Process* (MDP). In such a game (also called 1.5-player game), Player 2 plays according to a fixed distribution over successor states. Thus, each strategy of Player 1 gives rise to a stochastic process, i.e. to a distribution over infinite sequences of states. Given a goal, one may then ask which is the strategy of Player 1 which maximizes the probability of satisfying the goal. There are known algorithms for solving this problem for various classes of goals [BdA95].

¹ Player 1 states are represented by circles and Player 2 states by squares.

The notion of optimal strategy suggests a first solution to our original problem of dealing with losing states in a 2-player game. We can assume that Player 2 plays uniformly at random and compute an optimal strategy for Player 1. Although this approach may be of interest to some cases, it is worth considering alternative criteria, which do not make as strong an assumption on the behavior of Player 2. One such criterion is derived from the classical game-theoretic notion of *dominance* [OR94]. Given two strategies σ and σ' of Player 1, we say that σ *dominates* σ' if σ is always at least as good as σ' , and better than σ' in at least one case. Dominance induces a strict partial order over strategies, whose maximal elements are called *admissible* (or *undominated*) strategies. In Section 3, we prove that a strategy is admissible if and only if it is simultaneously strongly winning and cooperatively strongly winning (i.e., strongly winning with the help of Player 2).

We claim that admissible strategies represent a convincing notion of best-effort strategy. Similarly to optimality in a 1.5-player game, this notion is goal-independent. Moreover, it does not make any assumption on the behavior of Player 2. Dominance can also be useful for multi-criteria optimization problems. Differently from stochastic optimality, dominance gives rise to a *partial* order over strategies, which leaves room for another goal to be pursued.

Memory. A useful measure for the complexity of a strategy consists in evaluating how much memory it needs regarding the history of the game. In the simplest case, a strategy requires no memory at all: its decisions are based solely on the current state of the game. Such strategies are called *positional* or *memoryless* [GZ05]. In other cases, a strategy may require the amount of memory that can be provided by a finite automaton (*finite memory*), or more [DJW97].

The memory measure of a strategy is particularly important for the applications that we target in this paper. Since we are interested in actually implementing strategies in hardware or software, the simplest the strategy, the easiest and most efficient it is to implement.

We devote Section 4 to studying the memory requirements for various types of “good” strategies. In particular, we prove that all goals that have positional winning strategies also have positional *strongly* winning strategies. We also prove that for *prefix-independent* positional goals, all positional winning strategies are automatically strongly winning. The latter property is not valid for all positional goals. Thus, new algorithms are still needed to compute strongly winning strategies.

The situation is different for admissible strategies: there are games with positional goals that have no positional admissible strategy. Again, prefix-independent goals are particularly well-behaved: those that admit positional winning strategies also admit positional admissible strategies, as we show by presenting a simple algorithm which computes positional admissible strategies for these goals.

2 Definitions

We treat games that are played by two players on a finite graph, for an infinite number of turns. The aim of the first player is to obtain an infinite trace belonging to a fixed set of winning traces, while the aim of the second player is the exact opposite. In the literature, such

games are termed *two-player*, *turn-based*, *qualitative* and *zero-sum*. The following definitions make this framework formal.

A *game* is a tuple $G = (S_1, S_2, \delta, C, F)$ such that: S_1 and S_2 are disjoint finite sets of states; let $S = S_1 \cup S_2$, we have that $\delta \subseteq S \times S$ is the transition relation and $C : S \rightarrow \mathbb{N}$ is the coloring function, where \mathbb{N} denotes the set of natural numbers including zero. Finally, $F \subseteq \mathbb{N}^\omega$ is the *goal*. With an abuse of notation, we extend the coloring function to paths in the game graph, with the obvious meaning. We assume that games are non-blocking, i.e. each state has at least one successor in δ . Formally, $G' = (S_1, S_2, \delta', C, F)$, where $\delta' \cap (S_1 \times S) = \delta \cap (S_1 \times S)$ and $\delta' \cap (S_2 \times S) \subseteq \delta \cap (S_2 \times S)$.

A (finite or infinite) path in G is a (finite or infinite) path in the directed graph $(S_1 \cup S_2, \delta)$. If a finite path ρ is a prefix of a finite or infinite path ρ' , we also say that ρ' *extends* ρ . We denote by $first(\rho)$ the first state of a path ρ and by $last(\rho)$ the last state of a finite path ρ .

Strategies. A *strategy* is a function $\sigma : S^* \rightarrow S$ such that for all $\rho \in S^*$, $(last(\rho), \sigma(\rho)) \in \delta$. Our strategies are deterministic, or, in game-theoretic terms, *pure*. We denote Σ_G the set of all strategies in G . We do not distinguish *a priori* between strategies of Player 1 and Player 2. However, for sake of clarity, we write σ for a strategy that should intuitively be interpreted as belonging to Player 1, and τ for the (rare) occasions when a strategy of Player 2 is needed.

Consider two strategies σ and τ , and a finite path ρ , and let $n = |\rho|$. We denote by $Outc_G(\rho, \sigma, \tau)$ the unique infinite path $s_0 s_1 \dots$ such that (i) $s_0 s_1 \dots s_{n-1} = \rho$, and (ii) for all $i \geq n$, $s_i = \sigma(s_0 \dots s_{i-1})$ if $s_{i-1} \in S_1$ and $s_i = \tau(s_0 \dots s_{i-1})$ otherwise. We set $Outc_G(\rho, \sigma) = \bigcup_{\tau \in \Sigma_G} Outc_G(\rho, \sigma, \tau)$. For all $s \in S$ and $\rho \in Outc_G(s, \sigma)$, we say that ρ is *consistent* with σ . Similarly, we say that $Outc_G(s, \sigma, \tau)$ is consistent with σ and τ . We extend the definition of consistency from infinite paths to finite paths in the obvious way.

A strategy σ is *positional* (or *memoryless*) if $\sigma(\rho)$ only depends on the last state of ρ . Formally, for all $\rho, \rho' \in S^*$, if $last(\rho) = last(\rho')$ then $\sigma(\rho) = \sigma(\rho')$.

Dominance. Given two strategies σ and τ , and a finite path ρ , we set $val_G(\rho, \sigma, \tau) = 1$ if $C(Outc_G(\rho, \sigma, \tau)) \in F$, and $val_G(\rho, \sigma, \tau) = 0$ otherwise. Given two strategies σ and σ' , we say that σ' *dominates* σ if: (i) for all $\tau \in \Sigma_G$ and all $s \in S$, $val_G(s, \sigma', \tau) \geq val_G(s, \sigma, \tau)$, and (ii) there exist $\tau \in \Sigma_G$ and $s \in S$ such that $val_G(s, \sigma', \tau) > val_G(s, \sigma, \tau)$.

It is easy to check that dominance is an irreflexive, asymmetric and transitive relation. Thus, it is a *strict partial order* on strategies.

Uniform stochastic games. In this presentation, a uniform stochastic game (USG) is syntactically equivalent to a game. However, its semantics is different: in a USG $G = (S_1, S_2, \delta, C, F)$, in states in S_1 Player 1 chooses the successor state, while in S_2 the successor is chosen according to a uniform distribution over all successors. A USG is thus a Markov Decision Process. A state s and a strategy σ induce a stochastic process which is a Markov chain. We denote $\Pr_s^\sigma(X)$ the probability of event X in the Markov chain generated by σ that starts at s . For more information on stochastic processes, see [FV97].

Given a strategy σ and a state s , and assuming that F is a measurable subset of \mathbb{N}^ω , we denote $val_G(s, \sigma)$ the probability of winning using σ from state s , that is $\Pr_s^\sigma(\{\rho \in S^\omega \mid$

$C(\rho) \in F\}$). We say that a strategy σ^* is *optimal* if it maximizes the probability of winning, from all states. Formally, for all states s , we require $\text{val}_G(s, \sigma^*) = \sup_{\sigma \in \Sigma_G} \text{val}_G(s, \sigma)$.

2.1 Good Strategies

In the following, unless stated otherwise, we consider a fixed game $G = (S_1, S_2, \delta, C, F)$ and we omit the G subscript whenever the game is clear from the context.

Let ρ be a finite path in G , we say that a strategy σ is *winning from ρ* if, for all $\rho' \in \text{Outc}(\rho, \sigma)$, the color sequence of ρ' belongs to F . We say that ρ is *winning* if there is a strategy σ which is winning from ρ . The above definition extends to states, by considering them as length-1 paths. A state that is not winning is called *losing*.

Further, a strategy σ is *cooperatively winning from ρ* if there exists a strategy τ such that the color sequence of the unique infinite path $\text{Outc}(\rho, \sigma, \tau)$ belongs to F . We say that ρ is *cooperatively winning* if there is a strategy σ which is cooperatively winning from ρ . Intuitively, a path is cooperatively winning if the two players together can extend that path into an infinite path that satisfies the goal. Again, the above definitions extend to states, by considering them as length-1 paths.

We can now present the following set of winning criteria. Each of them is a possible definition of what a “good” strategy is.

- A strategy is *winning* if it is winning from all winning states. This criterion intuitively demands that strategies enforce victory whenever the initial state allows it.
- A strategy is *subgame perfect* if it is winning from all winning paths. This criterion states that a strategy should enforce victory whenever the current history of the game allows it.
- A strategy is *strongly winning* if it is winning from all winning paths that are consistent with it.
- A strategy is *cooperatively winning* (in short, *c-winning*) if it is cooperatively winning from all cooperatively winning states. This criterion essentially asks a strategy to be winning with the help of Player 2.
- A strategy is *cooperatively subgame perfect* (in short, *c-perfect*) if it is cooperatively winning from all cooperatively winning paths.
- A strategy is *cooperatively strongly winning* (in short, *cs-winning*) if it is cooperatively winning from all cooperatively winning paths that are consistent with it.
- A strategy is *admissible* if there is no strategy that dominates it. This criterion favors strategies that are maximal w.r.t. the partial order defined by dominance.
- A strategy is *optimal* if it is so in the corresponding uniform stochastic game. This criterion endorses strategies that are optimal, assuming that the adversary chooses her moves according to a uniform distribution.

The notions of winning, optimal and cooperatively winning strategies are customary to computer scientists [Tho95,AHK97]. The notion of subgame perfect strategy comes from classical game theory [OR94]. The introduction of the notion of strongly winning strategy is motivated by the fact that in the target applications game histories that are inconsistent with the strategy of Player 1 cannot occur. Being strongly winning is strictly weaker than being subgame perfect. In addition, there are games for which there is a positional strongly

winning strategy, but no positional subgame perfect strategy (see Figure 8 in the Appendix). The term “strongly winning” seems appropriate since this notion is a natural strengthening of the notion of winning strategy.

We say that a goal F is *positional* if for all games G with goal F , there is a positional strategy that is winning in G . Recently, necessary and sufficient conditions for a goal to be positional were identified [GZ05].

3 Comparing Winning Criteria

In this section, we compare the winning criteria presented in Section 2.1, taking as the main reference the definition of winning strategy. Figure 2 summarizes the relationships between the winning criteria under consideration. We start by stating the following basic properties.

Lemma 1. *The following properties hold:*

1. *all strongly winning strategies are winning, but not vice versa;*
2. *all subgame perfect strategies are strongly winning, but not vice versa;*
3. *all cs-winning strategies are c-winning, but not vice versa;*
4. *all c-perfect strategies are cs-winning, but not vice versa;*
5. *all games have a winning (respectively, strongly winning, subgame perfect, c-winning, cs-winning, c-perfect, optimal, admissible) strategy.*

Proof. The containments stated in (1) and (2) are obvious by definition. The fact that those containments are strict is easily proved by the game in Figure 6 in the Appendix. Similarly, statements (3) and (4) follow from the definitions and from the example in Figure 7 in the Appendix.

Regarding statement (5), the existence of a winning (respectively, strongly winning, subgame perfect, c-winning, cs-winning, c-perfect, optimal) strategy is obvious by definition. The existence of an admissible strategy can be derived from Theorem 11 from [Ber07]. ■

3.1 Strongly Winning Strategies

It is easy to check that winning strategies need not be strongly winning. Here, we give a sufficient condition that a goal can satisfy to ensure that all winning strategies are strongly winning.

A goal F is *shrinkable* iff for all $c\rho \in F$, with $c \in \mathbb{N}$ and $\rho \in \mathbb{N}^\omega$, $\rho \in F$. A goal F is *extensible* iff for all $\rho \in F$ and all $c \in \mathbb{N}$, $c\rho \in F$. A goal F is *prefix-independent* iff it is both shrinkable and extensible. Examples of common prefix-independent goals include Büchi, co-Büchi and parity goals.

Lemma 2. *If a goal F is shrinkable, then, for all games with goal F , all winning paths end in a winning state, and all c-winning paths end in a c-winning state.*

Proof. We prove the statement for winning paths, as the one regarding c-winning paths can be proved along similar lines. Let G be a game with a shrinkable goal F , let $\rho = s_0 \dots s_n$ be

a winning path and let σ be a strategy which is winning from ρ . Consider all infinite paths that extend ρ and are consistent with σ . These paths all satisfy the goal F . If we remove the prefix ρ from these paths, they still all satisfy F , since F is shrinkable. Consider the strategy σ' defined by: for all $\pi \in S^*$,

$$\sigma'(\pi) = \begin{cases} \sigma(s_0 \dots s_{n-1} \pi) & \text{if } \text{first}(\pi) = s_n, \\ \text{arbitrarily defined} & \text{otherwise.} \end{cases}$$

It is immediate that σ' is winning from s_n and therefore s_n is a winning state. ■

The following corollary states that if a goal is shrinkable, winning strategies confine the game in the winning region.

Corollary 1. *If a goal is shrinkable, for all winning strategies σ , and for all finite paths ρ consistent with σ , if $\text{first}(\rho)$ is winning then $\text{last}(\rho)$ is winning.*

Theorem 1. *If a goal F is prefix-independent, then, for all games with goal F , all positional winning strategies are strongly winning, and all positional c-winning strategies are cs-winning.*

Proof. We prove the statement for positional winning strategies, as the one regarding positional c-winning strategies can be proved along similar lines. Let G be a game with goal F , and let σ be a winning strategy for G . Let $\rho = s_0 \dots s_n$ be a winning path which is consistent with σ . By Lemma 2, s_n is a winning state. Let $\rho' = s_0 \dots s_n s_{n+1} \dots$ be an infinite path which extends ρ and is consistent with σ . Since σ is positional and winning from s_n , $C(s_n s_{n+1} \dots) \in F$. Since F is extensible, $C(\rho') \in F$. Therefore, σ is winning from ρ . Being ρ generic, we conclude that σ is strongly winning. ■

The positionality assumption is necessary in the above result. For a prefix-independent goal, it is easy to devise winning strategies that are not positional and not strongly winning. On the other hand, being prefix-independent is not necessary for ensuring that all positional winning strategies are strongly winning. For instance, safety and reachability goals are not prefix-independent, but they ensure said property. Finally, simple examples show that neither shrinkability nor extensibility alone can replace prefix-independence in the assumptions of the above result.

3.2 Admissible Strategies

In this section we provide a characterization of admissibility in terms of the simpler criteria of strongly winning and cooperatively strongly winning. Such characterization will be useful to derive further properties of admissible strategies. We start with a trivial property of winning paths.

Lemma 3. *If $s_0 s_1 \dots s_n$ is a winning path and $s_{n-1} \in S_1$, then $s_0 s_1 \dots s_{n-1}$ is a winning path.*

We can now prove the main theorem of this section.




Fig. 2: Comparing winning criteria.

Theorem 2. *A strategy is admissible if and only if it is strongly winning and cooperatively strongly winning.*

Proof. For the “if” part, let σ^* be a strategy which is both strongly winning and cs-winning. Assume that there is a strategy σ that is better than σ^* in at least one case. In particular, let s be a state and τ be a strategy of Player 2 such that $val(s, \sigma^*, \tau) = 0$ and $val(s, \sigma, \tau) = 1$. We can build a Player 2 strategy τ' such that $val(s, \sigma^*, \tau') = 1$ and $val(s, \sigma, \tau') = 0$, thus proving that σ does not dominate σ^* .

Let αu be the longest prefix common to both $Outc(s, \sigma, \tau)$ and $Outc(s, \sigma^*, \tau)$. Precisely, let $Outc(s, \sigma, \tau) = \alpha uv \dots$ and $Outc(s, \sigma^*, \tau) = \alpha uv^* \dots$, where $\alpha \in S^*$, $u, v, v^* \in S$, and $v \neq v^*$. Clearly, it must be $u \in S_1$. We have that αu is not a winning path because σ^* is strongly winning and still $val(s, \sigma^*, \tau) = 0$. By Lemma 3, αuv is not a winning path either. We also have that αu is c-winning, because $val(s, \sigma, \tau) = 1$. Since σ^* is cs-winning, αuv^* is also c-winning.

Then, we define the Player 2 strategy τ' as follows. We let τ' coincide with τ on all finite paths that are prefixes of αu . On all paths that are extensions of αuv , we let τ' behave in such a way to ensure that $val(s, \sigma, \tau') = 0$. This is possible because αuv is not a winning path. On all paths that are extensions of αuv^* , we let τ' behave in such a way to ensure that $val(s, \sigma^*, \tau') = 1$. This is possible because αuv^* is a c-winning path and σ^* is cs-winning. We conclude that σ does not dominate σ^* . Therefore, no strategy dominates σ^* .

Next, we prove the “only if” part. Let σ^* be an admissible strategy. By contradiction, assume that σ^* is not strongly winning. Therefore, there exist a winning path ρ and a strategy τ^* such that ρ is consistent with σ^* and τ^* and $val(\rho, \sigma^*, \tau^*) = 0$. Now, let σ be a strategy that is winning from ρ . Define another Player 1 strategy σ' as follows, for all $\pi \in S^*$:

$$\sigma'(\pi) = \begin{cases} \sigma(\pi) & \text{if } \pi \text{ extends } \rho, \\ \sigma^*(\pi) & \text{otherwise.} \end{cases}$$

We show that σ' dominates σ^* , which is a contradiction. Take any Player 2 strategy τ . If σ' and τ together do not give rise to the finite path ρ , σ' behaves exactly like σ^* . If σ' and τ together

do give rise to the path ρ , from that point on σ' behaves like σ , and therefore ensures victory. This proves that σ' always performs at least as well as σ^* . Finally, there is a case where σ' performs better than σ^* : we have $val(first(\rho), \sigma', \tau^*) = 1$ and $val(first(\rho), \sigma^*, \tau^*) = 0$, which concludes the proof of the contradiction.

Next, we show that σ^* is cs-winning. Let ρ be a c-winning path and assume by contradiction that, for all strategies τ , $val(\rho, \sigma^*, \tau) = 0$. Let $\sigma^\bullet, \tau^\bullet$ be a pair of strategies such that $val(\rho, \sigma^\bullet, \tau^\bullet) = 1$. Let σ be a strategy that behaves like σ^* , except that for all paths extending ρ it behaves like σ^\bullet . We prove that σ dominates σ^* . If the path ρ is not formed during the game, σ behaves exactly like σ^* . If the path ρ is formed, σ^* loses with certainty, while σ wins in at least one case, namely against τ^\bullet . This proves that σ dominates σ^* , which is a contradiction. ■

3.3 Optimal Strategies

A natural notion of good strategy in a losing state is given by optimal strategies in the corresponding uniform stochastic game (USG). In this section, we explore the relationship between optimality and the other winning criteria.

First of all, a trivial example (see Figure 5 in the Appendix) shows that neither winning strategies nor strongly winning strategies nor even admissible strategies need be optimal. The game in Figure 3 shows that the converse also holds, i.e. optimal strategies need not be




Fig. 3: A game showing that optimal strategies need not be winning.

winning (and therefore neither strongly winning nor admissible). Suppose that the goal is to reach state s_2 . In the corresponding USG, both strategies of Player 1 are optimal, yielding a probability 1 of winning. However, only the strategy that goes directly to s_2 is winning. So, in this case winning proves to be finer (more discriminating) than optimal.

Finally, one might wonder whether starting from losing states optimal strategies are always admissible. The game in Figure 1, already presented in the introduction, settles this question in the negative. Suppose the goal is to visit infinitely often s_0 . All states and all paths in the game are losing. In the corresponding USG, all strategies have value 0. However, the strategy that always picks s_1 dominates all others, since it has a chance of winning (namely, against the strategy that always picks s_0).

4 Memory

In this section, we study the amount of memory required by “good” strategies for achieving different kinds of goals. We are particularly interested in identifying those goals which admit

positional good strategies, because positional strategies are the easiest to implement. The problem of identifying goals that admit positional winning strategies and positional optimal strategies is well studied in the literature. Here, we focus on the winning criteria of “strongly winning” and “admissible”.

4.1 Positional Strongly Winning Strategies

For a game $G = (S_1, S_2, \delta, C, F)$ and a path $\rho = s_0 \dots s_n$ in G , define $\text{detach}(G, \rho)$ as the game obtained from G by adding a copy of the path ρ to it as a chain of new states ending in the original state s_n . Formally, $\text{detach}(G, \rho) = (S_1, S'_2, \delta', C', F)$, where $S'_2 = S_2 \cup \{s'_0, s'_1, \dots, s'_{n-1}\}$ and $s'_0, s'_1, \dots, s'_{n-1}$ are new distinct states not belonging to S_2 nor to S_1 . Then, $(s, t) \in \delta'$ iff either (i) $(s, t) \in \delta$, or (ii) $s = s'_i$ and $t = s'_{i+1}$, or (iii) $s = s'_{n-1}$ and $t = s_n$. Finally, the color labeling is defined by:

$$C'(s) = \begin{cases} C(s_i) & \text{if } s = s'_i \text{ for some } i \in \{0, \dots, n-1\}, \\ C(s) & \text{otherwise.} \end{cases}$$

The detach operation adds some states to the game, but the new states are not reachable from any of the old states. As a consequence, if games G and $\text{detach}(G, \rho)$ start from an old state s , they are indistinguishable to both players. The following lemma formalizes this observation, by stating in particular that the detach operation preserves the winning property of paths.

Notice that in the following we commit a slight abuse in language by identifying positional strategies in G and in $\text{detach}(G, \rho)$. This is justified by the fact that the new states introduced in $\text{detach}(G, \rho)$ only have one successor, thus giving no more choices to either player.

Lemma 4. *For all strategies σ , and finite paths ρ, ρ' in G , σ is winning from ρ in G if and only if it is winning from ρ in $\text{detach}(G, \rho')$.*

The following lemma provides the very reason for introducing the detach operation in the first place. Consider a winning path in the original game G . In order for this path to occur in a play, we may need the collaboration of Player 2. On the other hand, detaching this path allows us to force its occurrence, with no help from Player 2. As a consequence, the initial state s'_0 of the detached path is a winning state in the detached game, as stated by the following lemma.

Lemma 5. *Let $\rho = s_0 \dots s_n$ be a path in G and let $G' = \text{detach}(G, \rho)$. Let s'_0 be the new state added to G' in correspondence to s_0 . If ρ is a winning path in G , then s'_0 is a winning state in G' .*

Proof. Let σ be a strategy that is winning from ρ in G , and let $\rho' = s'_0 \dots s'_{n-1} s_n$ be the detachment of ρ added to G' . Define a strategy σ' in G' as follows. For all finite paths π in G' , if π extends ρ' , i.e. $\pi = \rho' \pi'$, set $\sigma'(\pi) = \sigma(\rho \pi')$. For all other paths in G' , σ' chooses an arbitrary move available to Player 1. It is easy to check that the set of infinite paths in G which are consistent with σ and extend ρ is color-equivalent to the set of infinite paths in G' which are consistent with σ' and start in s'_0 . Therefore, σ' is winning from s'_0 in G' . ■

With the help of the previous lemmas, we are ready to prove the following.

Theorem 3. *A goal is positional if and only if it admits positional subgame perfect strategies.*

Proof. The “if” part being obvious by definition, assume that the goal is positional. Let G be a game and let W be the set of winning paths of G . W may be infinite but it is certainly countable. Consider any ordering of W into ρ_0, ρ_1, \dots . Consider the sequence of games $(G_i)_{i \geq 0}$ defined by $G_0 = G$ and $G_{i+1} = \text{detach}(G_i, \rho_i)$. Additionally, consider the sequence of strategies $(\sigma_i)_{i \geq 0}$ defined by: σ_0 is any positional winning strategy in G_0 , and

$$\sigma_{i+1} = \begin{cases} \sigma_i & \text{if } \sigma_i \text{ is winning in } G_{i+1}, \\ \text{any positional winning strategy in } G_{i+1} & \text{otherwise.} \end{cases}$$

We prove that the sequence $(\sigma_i)_{i \geq 0}$ converges to a strategy σ^* within a finite number of steps. Specifically, we show that once a strategy σ occurs in the sequence and then is replaced by another one, it will not occur at any later point in the sequence. This fact, together with the fact that the number of positional strategies is finite, leads to the convergence of the sequence. Assume by contradiction that there exist indices $0 \leq a < b$ such that $\sigma_a \neq \sigma_{a+1} \neq \sigma_b$ and $\sigma_a = \sigma_b$. Since $\sigma_a \neq \sigma_{a+1}$, σ_a is not winning in G_{a+1} . Therefore, by repeated application of Lemma 4, σ_a (or equivalently, σ_b) cannot be winning in G_b , which is a contradiction.

Next, we prove that σ^* is subgame perfect in G . Let $\rho = s_0 \dots s_n$ be a winning path in G . There is an index $j > 0$ such that $G_j = \text{detach}(G_{j-1}, \rho)$. Let s'_0 be the “copy” of s_0 added to G_j by the detach operation. Since ρ is a winning path in G , by Lemma 4 it is still winning in G_{j-1} , and by Lemma 5 s'_0 is a winning state in G_j . Therefore, σ_j is winning from s'_0 . Since σ^* is the ultimate value for the sequence of strategies, σ^* is also winning from s'_0 . Take an infinite path ρ' in G , which extends ρ and that after ρ is consistent with σ^* . If we replace in ρ' the prefix ρ with $s'_0 \dots s'_{n-1} s_n$, we obtain an infinite path ρ'' in G_j , which is consistent with σ^* (because σ^* is positional) and color-equivalent to ρ' . Since σ^* is winning from s'_0 , $C(\rho'') \in F$ and therefore $C(\rho') \in F$, which concludes the proof. ■

Notice, however, that there are games with non-positional goals that have a positional winning strategy but no positional subgame perfect strategy. An immediate consequence of the above theorem is the following.

Corollary 2. *A goal is positional if and only if it admits strongly winning strategies.*

4.2 Positional Admissible Strategies

Since all admissible strategies are winning, admissible strategies require at least as much memory as winning strategies. The following example shows that there are goals which admit positional winning strategies, but do not admit positional admissible strategies.

Example 1. Consider the goal which requires to visit at least twice color 1 (in LTL notation, $\varphi = \Diamond(1 \wedge X\Diamond 1)$). Such goal is positional, as it satisfies the criterion of [GZ05].

Consider the game in Figure 4, with the above goal φ . In the figure, the color of each state appears next to it or above it. One can easily check that states s_0 and s_3 are winning, while




Fig. 4: Positional goals need not admit positional admissible strategies.

all others are losing. The only choice for Player 1 occurs in s_1 , where he can choose between s_2 and s_4 . If the game starts in s_0 , in order to ensure victory, Player 1 must choose s_4 after s_1 . On the other hand, suppose that the game starts in s_1 . State s_1 is cooperatively winning, since both players can cooperate and achieve victory by following the path $s_1 s_2 s_3 s_4 s_5^\omega$. By Theorem 2 for a strategy to be admissible it has to be both winning and cooperatively winning. Thus, Player 1 must choose s_2 after s_1 , if the game started in s_1 itself. In conclusion, the choice of any admissible strategy from s_1 depends on the past history of the game.

In the following, we give a sufficient condition to ensure that a goal admits positional admissible strategies. We start with a simple algorithm, which we later show yields positional admissible strategies for many goals of interest.

Computing positional admissible strategies. Suppose that we are given a game G with a positional goal F , and that we have an algorithm for computing the set of winning states and a positional winning strategy for all games with goal F . Consider the following procedure.

1. Compute the set of winning states Win and a positional winning strategy σ for G .
2. Remove from G the edges of Player 1 which start in Win and do not belong to σ .
3. In the resulting game, compute and return a positional cooperatively winning strategy.

In general, this procedure may return a strategy that is neither winning nor cooperatively winning. However, in the following we show that indeed it returns a strategy that is both winning and cooperatively winning in many cases of interest.

As far as the complexity of the procedure is concerned, assuming the usual graph-like adjacency-list representation for games, we obtain the same asymptotical complexity as finding a positional winning strategy for F . In particular, step 3 can easily be performed by attributing all states to Player 1 and then running the algorithm for a positional winning strategy.

In the following, we prove that for prefix-independent positional goals the above procedure returns an admissible strategy.

Theorem 4. *If a goal is positional and prefix-independent, then it admits positional admissible strategies.*

Proof. Let G be a game with goal F , where F is positional and prefix-independent. Let Win be the set of winning state of G , and consider the application of the above procedure to G . Let σ_1 be the strategy computed at step 1 of the procedure, and σ_3 the output strategy. Clearly, σ_3 is positional. We prove that it is also admissible. According to Theorems 1 and 2, it is sufficient to prove that σ_3 is winning and cooperatively winning.

By Corollary 1, whenever the game starts in a winning state, σ_1 confines the game in Win . Since σ_3 coincides with σ_1 on Win , σ_3 is winning.

Next, let G_2 be the game built at Step 2 of the procedure. We know that σ_3 is cooperatively winning in G_2 . We now prove that σ_3 is cooperatively winning in G . Let s be a cooperatively winning state in G . We prove that s is cooperatively winning in G_2 as well. Let σ, τ be a pair of strategies in G such that $C(\text{Outc}(s, \sigma, \tau)) \in F$. We consider the following two cases: (1) the infinite path $\text{Outc}(s, \sigma, \tau)$ does not visit Win ; then, strategies σ and τ are valid in G_2 and thus s is cooperatively winning in G_2 ; (2) let ρ be the shortest prefix of $\text{Outc}(s, \sigma, \tau)$ which ends in Win ; then, consider a new strategy σ' which behaves as follows, for all $\pi \in S^*$:

$$\sigma'(\pi) = \begin{cases} \sigma_1(\pi) & \text{if } \pi \text{ extends } \rho, \\ \sigma(\pi) & \text{otherwise.} \end{cases}$$

Consider strategies σ' and τ playing together in G_2 . At first, σ' behaves like σ and so the path ρ is built. Then, σ' behaves like the winning strategy σ_1 . Eventually, the infinite path $\rho\rho'$ is built, where $C(\text{last}(\rho)\rho') \in F$. Since F is extensible, $C(\rho\rho') \in F$ and thus s is cooperatively winning in G_2 . ■

This result proves that common goals such as Büchi, co-Büchi, and parity all admit positional admissible strategies. However, prefix-independence is not necessary for admitting positional admissible strategy. Further, it is not necessary for the procedure to work either. In particular, it is easy to prove that the procedure also returns an admissible strategy for reachability and safety goals.

5 Conclusions and Future Work

We studied and compared several criteria for establishing what are the preferable strategies for Player 1 in a given game. Being winning, or rather strongly winning, intuitively seems to be a necessary prerequisite for any best-effort strategy. Accordingly, optimality against a random opponent should be discarded as it does not imply winning. However, for some applications it could be worth considering a strategy which is primarily strongly winning, and optimal from paths that are not winning. Admissibility, on the other hand, passes all our tests, including the one of being efficiently computable for goals of interest. On the other hand, in general admissibility incurs some memory cost.

This preliminary study leaves several open problems, including: characterizing goals having positional admissible strategies, finding algorithms for computing strongly winning and admissible strategies for goals that are not prefix-independent.

Aside from the goal-independent best-effort criteria studied in this paper, it is worth noticing that goal-specific criteria offer an entirely different range of possibilities, many of which remain unexplored in the literature.

Acknowledgments. The author would like to thank Hugo Gimbert for a fruitful conversation on the subject.

References

- AHK97. R. Alur, T.A. Henzinger, and O. Kupferman. Alternating-time temporal logic. In *Proc. 38th IEEE Symp. Found. of Comp. Sci.*, pages 100–109. IEEE Computer Society Press, 1997.
- ALW89. M. Abadi, L. Lamport, and P. Wolper. Realizable and unrealizable specifications of reactive systems. In *Proc. 16th Int. Colloq. Aut. Lang. Prog.*, volume 372 of *Lect. Notes in Comp. Sci.*, pages 1–17. Springer-Verlag, 1989.
- BBF07. S. Baseline, P.A. Bonatti, and M. Faella. On interoperable trust negotiation strategies. In *POLICY 07: 8th IEEE International Workshop on Policies for Distributed Systems and Networks*. IEEE Computer Society, 2007.
- BdA95. A. Bianco and L. de Alfaro. Model checking of probabilistic and nondeterministic systems. In *Found. of Software Tech. and Theor. Comp. Sci.*, volume 1026 of *Lect. Notes in Comp. Sci.*, pages 499–513. Springer-Verlag, 1995.
- Ber07. D. Berwanger. Admissibility in infinite games. In *STACS*, volume 4393 of *Lect. Notes in Comp. Sci.*, 2007.
- dAFMR05. L. de Alfaro, M. Faella, R. Majumdar, and V. Raman. Code aware resource management. In *EMSOFT 05: 5th Intl. ACM Conference on Embedded Software*, pages 191–202. ACM Press, 2005.
- dAH01. L. de Alfaro and T.A. Henzinger. Interface theories for component-based design. In *EMSOFT 01: 1st Intl. Workshop on Embedded Software*, volume 2211 of *Lect. Notes in Comp. Sci.*, pages 148–165. Springer-Verlag, 2001.
- DJW97. S. Dziembowski, M. Jurdziński, and I. Walukiewicz. How much memory is needed to win infinite games? In *Proc. 12th IEEE Symp. Logic in Comp. Sci.* IEEE Computer Society, 1997.
- EJ91. E.A. Emerson and C.S. Jutla. Tree automata, mu-calculus and determinacy (extended abstract). In *Proc. 32nd IEEE Symp. Found. of Comp. Sci.*, pages 368–377. IEEE Computer Society Press, 1991.
- FV97. J. Filar and K. Vrieze. *Competitive Markov Decision Processes*. Springer-Verlag, 1997.
- GZ05. H. Gimbert and W. Zielonka. Games where you can play optimally without any memory. In *CONCUR 05: Concurrency Theory. 16th Int. Conf.*, volume 3653 of *Lecture Notes in Computer Science*. Springer, 2005.
- KR01. S. Kremer and J.-F. Raskin. A game-based verification of non-repudiation and fair exchange protocols. In *Proceedings of the 12th International Conference on Concurrency Theory*, volume 2154 of *Lecture Notes in Computer Science*, pages 551–565. Springer, 2001.
- MPS95. O. Maler, A. Pnueli, and J. Sifakis. On the synthesis of discrete controllers for timed systems. In *Proc. of 12th Annual Symp. on Theor. Asp. of Comp. Sci.*, volume 900 of *Lect. Notes in Comp. Sci.* Springer-Verlag, 1995.
- OR94. M.J. Osborne and A. Rubinstein. *A Course in Game Theory*. MIT Press, 1994.
- PR89. A. Pnueli and R. Rosner. On the synthesis of a reactive module. In *Proceedings of the 16th Annual Symposium on Principles of Programming Languages*, pages 179–190. ACM Press, 1989.
- Tho95. W. Thomas. On the synthesis of strategies in infinite games. In *Proc. of 12th Annual Symp. on Theor. Asp. of Comp. Sci.*, volume 900 of *Lect. Notes in Comp. Sci.*, pages 1–13. Springer-Verlag, 1995.
- Zie98. W. Zielonka. Infinite games on finitely coloured graphs with applications to automata on infinite trees. *Theoretical Computer Science*, 200:135–183, June 1998.

Appendix: Additional Examples




Fig. 5: Admissible strategies need not be optimal. The goal is to reach s_4 . All strategies of Player 1 are winning, strongly winning and admissible. On the other hand, only the strategy that chooses s_1 is optimal, leading to a probability of winning of $\frac{1}{2}$.




Fig. 6: A game having a winning strategy which is not strongly winning, and a strongly winning strategy which is not subgame perfect. The goal is $\varphi = 1 \wedge \Diamond \Box 2$. The positional strategy consisting of going from s_1 to s_0 (the thick edge) is winning but not strongly winning. The strategy that chooses s_1 when the current history contains exactly one occurrence of s_0 , and chooses s_0 otherwise is strongly winning but not subgame perfect, due for instance to the winning path $s_0 s_1 s_0 s_1$.




Fig. 7: A game having a c-winning strategy which is not cs-winning, and a cs-winning strategy which is not c-perfect. The goal is $\varphi = 1 \wedge \Diamond \Box 2$. The only c-winning state is s_0 . The positional strategy $s_0 \rightarrow s_1, s_2 \rightarrow s_2$ is c-winning but not cs-winning. Consider the strategy that chooses $s_0 \rightarrow s_1$ and then (i) $s_2 \rightarrow s_1$ if s_1 was visited in the current history and (ii) $s_2 \rightarrow s_2$ if s_1 was not visited in the current history. Such strategy is cs-winning but not c-perfect, due to the c-winning path $s_0 s_2$.




Fig.8: A game having a positional strongly winning strategy, but no positional subgame perfect strategy. The goal is $\varphi = 0 \wedge ((\Diamond 1 \wedge \Diamond 3) \vee (\Diamond 2 \wedge \Diamond 4))$. The positional strategy $s_0 \rightarrow s_1, s_3 \rightarrow s_4$ is strongly winning. All subgame perfect strategies need memory in state s_3 .